pBG

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER OF PATENTS AND TRADEMARKS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/703,909	11/01/2000	Todd P. Guay	1958.1031-000 (OID-1999-0	3199

21005 7590 05/20/2003

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.
530 VIRGINIA ROAD
P.O. BOX 9133
CONCORD, MA 01742-9133

EXAMINER	
SCHRANTZ, STEPHEN D	
ART UNIT	**PAPER NUMBER**
2177	5

DATE MAILED: 05/20/2003

Please find below and/or attached an Office communication concerning this application or proceeding.

PRQ

Office Action Summary	Application No.	Applicant(s)
	09/703,909	GUAY ET AL.
	Examiner	Art Unit
	Steve Schrantz	2177

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).
- Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *04 March 2003*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-18 and 20-136* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-18 and 20-136* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

11)☐ The proposed drawing correction filed on _____ is: a)☐ approved b)☐ disapproved by the Examiner.

If approved, corrected drawings are required in reply to this Office action.

12)☐ The oath or declaration is objected to by the Examiner.

Priority under 35 U.S.C. §§ 119 and 120

13)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____ .

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

14)☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. § 119(e) (to a provisional application).

a) ☐ The translation of the foreign language provisional application has been received.

15)☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. §§ 120 and/or 121.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO-1449) Paper No(s) _____ .
4) ☐ Interview Summary (PTO-413) Paper No(s). _____ .
5) ☐ Notice of Informal Patent Application (PTO-152)
6) ☐ Other:

DETAILED ACTION

Response to Amendment

This office action [paper #5] is in response to Amendment [paper #4] filed 3/7/03.

In paper #4, the applicant amended claims 1, 5-8, 11, 12, 27, 32, 53, and 54. The applicant also added new claims 55-136.

Claim Rejections - 35 USC § 103

1. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

2. Claims 1-6, 8-12, 18, 20, 22-24, 26-30, 32, 34-38, and 44-46, 48-50, 52-54, 56, 58-60, 62-64, 66-70, 76-77, 79-81, 83-86, 88-90, 92-96, 102-103, 105-107, 109-113, 115-117, 119-123, 129-130, 132-134, and 136 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri et al. (U.S. Patent 5,926,813, now referred to as Chaudhuri (I)) and Chaudhuri et al. (U.S. Patent 6,223,171, now referred to as Chaudhuri (II)).

3. Chaudhuri (I) teaches independent claims 1, 53-54, and 110 by the following:
"forming an index superset from a union of a current index set and a proposed index set" at col. 7 lines 56-63. Chaudhuri I teaches the storage of what-if indexes [proposed indexes that are absent from the database server] and existing indexes [current index set]. The storage of the indexes serves as an index superset.

"deriving a candidate index set from the index superset, the derived candidate index set being included in the plurality of candidate index sets" at col. 7 line 64 to col. 8 line 10. Chaudhuri (I) teaches that the superset of proposed index sets and existing index sets are considered in evaluating index configurations [candidate index sets].

"generating statistics based on the derived candidate index set" at col. 10 lines 13-25. Chaudhuri (I) teaches that costs [generated statistics] for each candidate index configuration are determined and stored.

Chaudhuri (I) does not teach a user interface that presents the general statistics. Chaudhuri (II) presents the statistics at Fig. 6-8 and Fig. 10-15. It would have been obvious to one ordinarily skilled in the art at the time of the invention to present these statistics to the administrator or user because of the advantages of allowing the user to analyze the data. By presenting the statistics, the user can gain a better understanding of the costs and the advantages of particular indexes with each workload. Through the analysis, the user can determine the potential impact and improvements in performance of a candidate index set on a particular workload as taught at Chaudhuri II col. 1 line 64 to col. 2 line 13.

Claims 53-54 refers to "repeatedly derives a candidate index set" which can be found at Chaudhuri (I) col. 7 lines 46-56.

4. Chaudhuri (I) teaches dependent claims 2, 58, 84, and 111 by the following: "generating current index statistics for the workload responsive to the current index set, the presented generated statistics comprising the generated current index statistics" at col. 10 lines 13-21.

5. Chaudhuri (I) teaches dependent claims 3, 59, 85, and 112 by the following:

"repeatedly deriving a candidate index set and generating statistics based on the proposed index set" at col. 7 lines 46-56.

6. Chaudhuri (I) and Chaudhuri (II) teach dependent claims 4, 60, 86, and 113 by the following:

"terminating the repeated execution when at least one candidate index solution is found that adheres to user-imposed constraints" at Chaudhuri (I) col. 6 lines 9-35 and the user imposed constraints can be found at Chaudhuri (II) Fig. 4. The user is able to choose those indexes that can be removed and those indexes will not degrade the performance found at col.11 line 66 to col. 12 line 6;

"... no further indexes can be removed from said candidate index solution without degrading performance of the workload" at Chaudhuri (I) col. 19 lines 48-58. Chaudhuri (I) teaches a greedy algorithm that selects the indexes to be used to reduce the cost of the workload. If these indexes will not reduce the cost, the index will not be added to the set. Chaudhuri (II) in Fig. 4 also demonstrates a user interface that allows an administrator to add and remove indexes in order to achieve improved performance;

"... without disabling an integrity constraint" at Chaudhuri (I) col. 1 lines 52-55.

7. Chaudhuri (I) teaches dependent claims 5, 63, 89, and 116 by the following:

"wherein deriving the baseline statistics comprises disabling current indexes" at col. 14 lines 60-65.

8. Chaudhuri (I) teaches dependent claims 6, 64, 90, and 117 by the following: "wherein generating statistics for a statement comprises generating at least one statistic based on an execution plan created by an optimizer" at col. 10 lines 53-56 and col. 8 lines 4-7.

9. Chaudhuri (I) teaches dependent claims 8, 66, 92, and 119 by the following: "wherein the execution plan is based on statistics for at least one schema object accessed by the statement" at col. 10 lines 53-56 and col. 5 lines 52-56. The query optimizer may user an index over any single table of the database. The costs include the index configuration of both the indexes and the tables. The statistics are based upon the schema objects.

10. Chaudhuri (I) teaches dependent claims 9, 35, 67, 93, and 120 by the following: "the at least one schema object is a table" at col. 12 lines 20-29.

11. Chaudhuri (I) teaches dependent claims 10, 36, 68, 94, and 121 by the following: "the at least one schema object is an index" at col. 10 lines 54-55.

12. Chaudhuri (I) teaches dependent claims 11, 69, 95, and 122 by the following: "for a table accessed by a statement under evaluation, using the execution plan to identify at least one index that would be used to retrieve data from the table upon an execution of the statement" at col. 6 lines 9-13 and col. 10 lines 53-56. The execution plan is returned with a cost estimate of executing a designated query for the designated candidate index configuration. The index

selection tool attempts to select an index configuration that is optimal, so the cost will lead to the identification of an index.

13. Chaudhuri (I) teaches dependent claims 12, 70, 96, and 123 by the following: "the optimizer generates a cost of the execution plan" at col. 10 lines 48-56. The query optimizer returns an execution plan with a cost estimate. Because the optimizer returns both the execution plan and the cost estimate, the optimizer is considered to have generated both the plan and the cost.

14. Chaudhuri (I) teaches dependent claims 18, 44, 76, 102, and 129 by the following: "wherein the statistics include an index usage" at col. 10 lines 38-47 and col. 10 lines 53-56.

15. Chaudhuri (I) teaches dependent claims 20, 46, 77, 103, and 130 by the following: "wherein the statements are SQL statements" at col. 5 lines 40-41 and col. 5 lines 61-64.

16. Chaudhuri (I) teaches dependent claims 22, 79, 105, and 132 by the following: "deriving a candidate index set is responsive to a predetermined maximum number of allowed indexes" at col. 13 lines 1-9.

17. Chaudhuri (II) teaches dependent claims 23, 49, 80, 106, and 133 by the following: "wherein deriving a candidate index set is responsive to available storage space" at col. 1 lines 24-34.

18. Chaudhuri (II) teaches dependent claims 24, 50, 81, 107, and 134 by the following:

"the proposed index set is provided by a user" at col. 12 lines 1-10.

19. Chaudhuri (I) teaches dependent claims 26, 52, 83, 109, and 136 by the following:

"an execution plan is created without creating indexes which are not in the current index set" at col. 7 lines 52-63.

20. Chaudhuri (I) teaches dependent claim 27 by the following:

"a workload evaluator which evaluates each statement within the workload" at col. 10 lines 22-25. Chaudhuri I teaches that each statement [query] in the workload is evaluated.

"an index solution evaluator which, responsive to the workload evaluator, evaluates each index in a candidate index set with respect to the workload, the candidate index solution being one of the plurality of candidate index sets" at col. 10 lines 16-18. Chaudhuri (I) teaches that each workload is evaluated against each candidate index set.

"each candidate index set derived from an index superset formed by the union of a current index set and a proposed index set" at col. 7 line 64 to col. 8 line 10. Chaudhuri (I) teaches that the superset of proposed index sets and existing index sets are considered in evaluating index configurations [candidate index sets].

"a solution/rollup evaluator which, responsive to the index solution evaluator, evaluates the candidate index solution" at col. 10 lines 38-47;

"a solution refiner which, responsive to the solution/rollup evaluator, generates at least one new candidate index solution" at col. 2 lines 59-62.

21. Chaudhuri (II) teaches dependent claim 28 by the following:

"the solution refiner generates at least one new candidate index solution by eliminating at least one index within the candidate index solution that does not adhere to user-imposed constraints" at Fig. 4. The solution refiner allows the user to create a new candidate solution by removing existing indexes. Because the user selects the indexes to be removed, the indexes do not adhere to a user-imposed constraint.

22. Chaudhuri (II) teaches dependent claim 29 by the following:

"wherein the constraint is a user-defined constraint" at col. 11 line 46 to col. 12 line 30. This segment teaches that an administrator can define many aspects, which will be later used to determine the proposed index. The administrator can set the workload, the index configuration, and the configurations for the database system. The segment teaches many of the administrative functions of the database allowed to the user. Later, the statistical information acquired from the administrator's defined constraints (costs determined from the workload at col. 15 lines 18-26) will limit the indexes (col. 15 lines 6-11 and col. 15 lines 58-63).

23. Chaudhuri (II) teaches dependent claim 30 by the following:

"the constraint is a memory-usage constraint" at col. 1 lines 24-34.

24. Chaudhuri (I) teaches dependent claim 32 by the following:

"wherein the workload evaluator evaluates an execution plan created by an optimizer" at col. 10 lines 48-56 Chaudhuri (I) teaches that the execution plan is created by an optimizer and returns the cost of the plan also. Furthermore, Chaudhuri (I) teaches the workload evaluator [cost evaluation tool] that evaluates every query in the workload.

"the execution plan comprising, for each statement of the workload, an execution plan which represents a series of steps for executing the statement" at Chaudhuri I col. 10 lines 53-56. The query optimizer creates an execution plan for the query [statement]. It is well known that the execution plan involves the steps for executing a statement. Chaudhuri II teaches that the execution plan consists of the indexes to use, the operators, and other information at col. 17 lines 34-44.

"the workload evaluator further generating and recording statistics based on the evaluation of the execution plan" at col. 10 lines 48-56.

25. Chaudhuri (I) teaches dependent claim 34 by the following:

"creating an execution plan is based on statistics for at least one schema object accessed by the statement" at col. 10 lines 53-56 and col. 5 lines 52-56.

26. Chaudhuri (I) teaches dependent claim 37 by the following:

"wherein the workload evaluator, for a table accessed by a statement under evaluation, identifies at least one index which would be used to retrieve data from the table upon an execution of the statement" at col. 18 lines 40-48.

27. Chaudhuri (I) teaches dependent claims 38 and 45 by the following:

"determining a cost of the execution plan" at col. 10 lines 54-56.

28. Chaudhuri (I) teaches dependent claims 48 by the following:

"deriving a candidate index set is responsive to a predetermined maximum number of allowed indexes" at col. 13 lines 1-9.

29. Chaudhuri (I) teaches dependents claims 56, 62, 88, and 115 by the following:

"generating baseline statistics for each statement in the workload, wherein generating statistics is additionally based on the baseline statistics" at col. 6 lines 14-22 and col. 10 lines 48-56. Chaudhuri (I) teaches that the query optimizer estimates a cost for each query [statement] in the workload. The estimates for each query are then summed to determine statistics of the workload with the candidate index set.

30. Claims 7, 33, 65, 91, and 118 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri (I) et al. (U.S. Patent 5,926,813) and Chaudhuri et al. (II) (U.S. Patent

6,223,171) as applied to claim 1 above, and further in view of Jakobsson et al. (U.S. Patent 5,924,088).

As per claims 7, 33, 65, 91, and 118, Chaudhuri (I) teaches that an execution plan is created at col. 10 lines 53-56. Chaudhuri (I) also teaches that the indexes are used in order to access the database at col. 3 lines 50-54. Chaudhuri I and II do not teach that the "execution plan is based on available access paths". Jakobsson does teach the use of access paths at col. 4 lines 1-4. It would have been obvious to one ordinarily skilled in the art at the time of the invention to base the execution plan on an access path because the performance characteristics for a data retrieval may vary greatly depending on the choice of index access path as taught at col. 4 lines 8-12. The costs of the statements of the workload can be calculated through the use of the access paths. Using these paths would allow for improved calculations of the costs of each statement.

31. Claims 13-16, 39-41, 71-74, 97-100, and 124-127 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri et al. (I) (U.S. Patent 5,926,813) and Chaudhuri et al. (II) U.S. Patent (6,223,171) as applied to claim 1 above, and further in view of Eberhard et al. (U.S. Patent 6,003,022).

32. As per claims 13, 39, 71, 97, and 124, Chaudhuri (I) describes an invention that selects an index that would work best for a workload as described above. Chaudhuri (I) also teaches that each query does have a particular cost associated with each particular index at col. 10 lines 14-19. Chaudhuri (I and II) do not teach that the cost of the execution plan is derived from a resource. Eberhard does teach that the execution costs are derived from a resource at col. 3 lines

43-45. It would have been obvious to on ordinarily skilled in the art at the time of the invention to derive the cost of the execution plan from a particular resource because of execution costs in both CPU and I/O as taught at col. 3 lines 43-59. By calculating the costs of the queries under a particular index, the invention would be capable of calculating how each candidate index can affect each statement in the workload. Deriving the cost of the execution plan from a resource use needed to execute the statement allows the system to know how each query will affect the system's performance. Considering the system's performance will allow a better cost to be derived for the indexing system.

33. Eberhard teaches dependent claims 14, 40, 72, 98, and 125 by the following: "the resource use includes CPU execution time" at col. 3 lines 43-44.

34. Eberhard teaches dependent claims 15, 41, 73, 99, and 126 by the following: "the resource use includes input/output access" at col. 3 lines 43-44.

35. Eberhard teaches dependent claims 16, 42, 74, 100, and 127 by the following: "the statistics include the number of executions of the statement" at col. 24 lines 51-55.

36. Claims 17, 43, 75, 101, and 128 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri et al. (I) (U.S. Patent 5,926,813) and Chaudhuri et al. (II) (U.S. Patent 6,223,171) as applied to claim 1 above, and further in view of Smith et al. (U.S. Patent 5,404,510).

37. As per claims 17, 43, 75, 101, and 128, Chaudhuri (I) and Chaudhuri (II) teaches a query optimizer which searches for the best index to evaluate a given workload. Both Chaudhuri I and II do not teach that a user is capable of defining the importance of the statement. Smith does teach a user-defined importance of the statement at col. 7 lines 22-32. It would have been obvious to one ordinarily skilled in the art at the time of the invention to allow a user to define the importance of the statements found in the workload because the index selection tool could evaluate the statistics of the indexes according to the importance of the statements. A user could rank the importance of the queries in the workload. The ranking could then be used to define an index that works best for the workload. The index would then be created in consideration of the higher ranked statements.

38. Claims 21, 47, 78, 104, and 131 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri (I) et al. (U.S. Patent 5,926,813) and Chaudhuri et al. (II) (U.S. Patent 6,223,171) as applied to claim 1 above, and further in view of Celis et al. (U.S. Patent 6,021,405).

39. As per claims 21, 47, 78, 104, and 131, Chaudhuri (I) (5,926,813) teaches the use of statistics to find the cost of statements using a particular index. Both Chaudhuri I and II do not teach that the workload is reduced into unique statements in order to determine these costs. Celis teaches that the workload is reduced into unique expressions at col. 5 lines 37-44.

Chaudhuri I teaches that the index selection tool is expensive and that each request to run query optimizer increases the cost of time and memory at col. 6 lines 23-35. Furthermore, Chaudhuri I teaches that the queries of a workload are sent to the query optimizer at col. 10 lines 13-60. Chaudhuri I also teaches that the cost evaluation tool attempts to reduce the number of invocations of query optimizer by determining costs of queries of workloads without invoking query optimizer at col. 10 lines 61-67. By not sending particular queries to the query optimizer, the workload is reduced. In particular, Chaudhuri I teaches the elimination of running repetitions of the query on atomic indexes. By not running repetitions of a query on similar indexes, Chaudhuri I is reducing the workload into unique statements. Celis teaches in greater depth that redundant expressions are removed from the queries in order to save cost as taught at col. 1 lines 38-51.

It would have been obvious to one ordinarily skilled in the art at the time of the invention to generate the statistic using only unique expressions because of the ability to reduce processing time for redundant queries as taught at Celis col. 1 lines 41-51. Celis's invention is a query optimizer that is used to reduce the cost of each statement as taught at col. 1 lines 26-31. By only optimizing the unique statements, the optimizer will not be burdened by redundant statements. Chaudhuri I teaches the reduction of the number of invocations of query optimizer at col. 6 lines 51-59. By reducing the workload into unique statements, Chaudhuri I will invoke query optimizer fewer times for statements that had previously been optimized. The needless optimizations cost more processing time and ultimately increase the execution time for the query as taught at Celis col. 1 lines 42-51 and Chaudhuri I col. 6 lines 51-59.

40. Claims 25, 51, 82, 108, and 135 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri et al. (I) (U.S. Patent 5,926,813) and Chaudhuri et al. (II) (U.S. Patent 6,223,171) as applied to claim 1 above, and further in view of Lomet et al. ("The Microsoft Database Research Group").

Chaudhuri (I) teaches a rules based system that uses particular functions to find an index configuration for a given workload at col. 5 line 57 to col. 6 line 13. Both Chaudhuri I and II do not specifically teach an expert system that chooses the index configuration. Lomet does teach an expert system that chooses a proposed index at pg. 83. It would have been obvious to one ordinarily skilled in the art at the time of the invention to use an expert system because of the ability of expert systems to make index selections. By using an expert system, the administrator would not have to determine the most useful index set. The expert system could use prior knowledge and the calculations found in Chaudhuri (I) to determine an index set.

41. Chaudhuri (I) teaches dependent claims 25, 51, 82, 108, and 135 by the following: "wherein the proposed index set is provided by an expert system" at col. 5 line 57 to col. 6 line 13.

42. Claim 31 is rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri et al. (I) (U.S. Patent 5,926,813) and Chaudhuri et al. (II) (U.S. Patent 6,223,171) and Jakobsson et al. (U.S. Patent 5,924,088) as applied to claim 7 above, and further in view of Siegel et al. ("A Method for Automatic Rule Derivation to Support Semantic Query Optimization").

43. Jakobsson teaches dependent claim 31 by the following:

"the solution refiner generates at least one new candidate index solution by eliminating at least one index on a small table under evaluation, and wherein the at least one index does not enforce an integrity constraint" at col. 12 line 66 to col. 13 line 1.

Jakobsson teaches that the cost of searching a table is directly related with the size of the table. Jakobsson teaches that indexes can be used on these smaller tables. Jakobsson also teaches the use of different indexes to use in the index access path at col. 13 line 17-25 and Figs. 4-5. Neither Jakobsson nor Chaudhuri (I and II) teach that the indexes removed do not enforce an integrity constraint. Siegel does teach that indexes can be a part of an integrity constraint at pg. 1.

It would have been obvious to one ordinarily skilled in the art at the time of the invention to not remove indexes that were a part of the integrity constraint because of the problems that can arise by removing integrity constraints. The indexes ensure the validity of the database at pg. 1. By removing integrity constraints, the database would no longer be guaranteed to contain valid information. These constraints can also be used to decrease the retrieval time of a statement as taught at pg. 1.

44. Claims 55, 57, 61, 87, and 114 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri et al. I (U.S. Patent 5,926,813) and Chaudhuri et al. II (U.S. Patent 6,223,171) as applied to claim 1 above, and further in view of Gurry et al. (Oracle Performance Tuning).

45. As per claim 55, 57, 61, 87, and 114, Chaudhuri I teaches the estimation of queries over candidate index sets as taught at col. 10 lines 48-56. Chaudhuri I and II do not teach that the index volatility is included in the statistics. Gurry does teach the index volatility in the statistics at pages 353-354. Gurry teaches the volatility of indexes through the HEIGHT and DEL_LF_ROWS columns found in the statistics. If the statistics for the HEIGHT or DEL_LF_ROWS exceed a particular number, the index is a candidate for dropping and recreating. The dropping and recreating of an index is a part of the volatility of the index. It would have been obvious to one ordinarily skilled in the art at the time of the invention to include the index volatility in the statistics because of the costs incurred by dropping and recreating indexes. The index volatility demonstrates to the user the possibility of the index to be dropped and recreated. The volatile indexes may need to be rebuilt in order to obtain optimal performance [pg. 354]. Indexes over volatile tables may also become a space management problem [pg. 354]. The volatility statistics allow a DBA to be aware of the problems that may arise over particular index selections.

Response to Arguments

46. Applicant's arguments filed 3/4/03 have been fully considered but they are not persuasive (see examiner's response discussed below).

47. As per claims 1, 27, 53, and 54, the applicant states that Chaudhuri I starts with the set of all possible indexes and then attempts to reduce the number of indexes at the bottom of page 14 and top of page 15.

The examiner disagrees with the applicant in Chaudhuri's [I] formation of an index superset. At col. 7 lines 46-63, Chaudhuri (I) teaches a superset of indexes that consist of what-if (proposed) index entries and existing index entries. Chaudhuri (I) discusses gathering statistics for each proposed index and the existing index. Furthermore, Chaudhuri (I) teaches the creation of a superset by storing each index in a similar manner [the difference is that the proposed indexes are marked]. The superset created by Chaudhuri (I) consists of sets of both proposed [what-if] indexes and existing indexes.

The applicant (bottom of page 14 and top of page 15) states that Chaudhuri I starts with all possible indexes and then reduces the number of indexes and therefore index configuration for evaluation at col. 6 lines 62-67. By starting with all possible indexes, Chaudhuri (I) teaches a superset of all possible indexes, indexes that have existed and new ones that are proposed. Chaudhuri I does not teach that all of the indexes are evaluated. The candidate index selection tool determines a set of candidate indexes for evaluation as taught at col. 6 lines 62-67.

48. As per claims 6 and 32, the applicant argues that the references do not teach or suggest that the cost evaluation tool looks at and evaluates the [execution] plan itself at page 16 middle of the page.

The amended claim does not state that the execution plan is evaluated. At col. 10 lines 53-57, Chaudhuri (I) does teach that the query optimizer returns an execution plan that comprises the cost estimate, an evaluation of the execution plan.

49. As per claims 21 and 47 [argument found at the bottom of page 16], the applicant states that it would not be obvious to one ordinarily skilled in the art at the time of the invention to combine the Celis reference with Chaudhuri I and Chaudhuri II.

The examiner disagrees with the lack of motivation for combining the three references. Chaudhuri I teaches the optimization of queries associated with the cost evaluation as shown at Fig. 3 References 240, 241, 322, and 324. Celis better describes the optimization of queries as shown in the title.

In the middle of page 17, the applicant states that Celis does not teach that the workload is reduced into unique statements. The examiner disagrees with the applicant's argument.

Chaudhuri I teaches that the index selection tool is expensive and that each request to run query optimizer increases the cost of time and memory at col. 6 lines 23-35. Furthermore, Chaudhuri I teaches that the queries of a workload are sent to the query optimizer at col. 10 lines 13-60. Chaudhuri I also teaches that the cost evaluation tool attempts to reduce the number of invocations of query optimizer by determining costs of queries of workloads without invoking query optimizer at col. 10 lines 61-67. By not sending particular queries to the query optimizer, the workload is reduced. In particular, Chaudhuri I teaches the elimination of running repetitions of the query on atomic indexes. By not running repetitions of a query on similar indexes, Chaudhuri I is reducing the workload into unique statements. Celis teaches in greater depth that redundant expressions are removed from the queries in order to save cost as taught at col. 1 lines 38-51.

50. As per claims 25 and 51, the applicant argues that Lomet teaches away from an expert system at page 18.

The expert system may not be connected to the query optimizer, but the expert wizard can assist with the index selection as taught at Lomet page 83. Lomet teaches that the expert system does not address the fundamental issues satisfactorily, but it does address the issues of an index selection tool [maybe not satisfactorily but it does address the issues]. The expert system has been encoded with rules to assist the index selection tool. The index-tuning wizard also functions as an expert system in that it creates hypothetical indexes and evaluates their potential for performance enhancement with respect to the workload. It filters out spurious indexes and exploits characteristics of the relational query engine as taught at page 83. It also selects a set of appropriate indexes for a workload. The wizard simulates the judgment and behavior of a human by analyzing the indexes and selecting the appropriate indexes.

51. As per claim 31, the applicant states that the Jakobsson does not teach the removal of indexes on small tables at the bottom of page 18.

Jakobsson teaches that the costs are less expensive on small tables at col. 12 line 66 to col. 13 line 5. Because these costs are less expensive, it is more likely that fewer indexes are needed. Furthermore, Chaudhuri I teaches the removal of indexes in the candidate index set at col. 15 lines 5-38. The candidate index set can contain any number of indexes. The example shows a candidate set of one or two indexes {I1, I2} and {I1}. The index I2 has been removed from the candidate index on the query of one table, onektup. Chaudhuri II also teaches the removal of indexes at Fig. 4.

Conclusion

References Cited

Chaudhuri et al.	U.S. Patent 5,926,813
Chaudhuri et al.	U.S. Patent 6,223,171
Jakobsson et al.	U.S. Patent 5,924,088
Eberhard et al.	U.S. Patent 6,003,022
Smith et al.	U.S. Patent 5,404,510
Celis et al.	U.S. Patent 6,021,405
Lomet et al.	"The Microsoft Database Research Group"
Siegel et al.	"A Method for Automatic Rule Derivation to Support Semantic Query Optimization"
Gurry et al.	Oracle Performance Tuning

52. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37

CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

53. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Steve Schrantz whose telephone number is (703) 305-7690. The examiner can normally be reached on Mon-Fri. 8:30-5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John Breene can be reached on (703) 305-9790. The fax phone numbers for the organization where this application or proceeding is assigned are (703) 746-7239 for regular communications and (703) 746-7238 for After Final communications. The TC 2100 Customer Service number is (703) 306-5631.

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the receptionist whose telephone number is (703) 305-3900.

Steve Schrantz
May 14, 2003

SRIRAMA CHANNAVAJJALA
PRIMARY EXAMINER